

"Express Mail" mailing label no. EE796280895US

Date of Deposit: 26 APRIL 2001

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the

Box Patent Application
Commissioner for Patents
Washington, D.C. 20231.

Jeffrey Slusher
(Typed or printed name of person mailing paper or fee)

Jeffrey Slusher
Signature of person mailing paper or fee)

Atty. ref: VMware11

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Attached is a new U.S. Utility Patent Application for:

TITLE: Undefeatable Transformation for Virtual Machine I/O Operations

INVENTORS:

Name: Carl A. WALDSPURGER
Residence: Atherton, California, USA
Address: 27 Park Dr., Atherton, CA 94027
Citizen of: USA

Name: Edouard BUGNION
Residence: Palo Alto, California, USA
Address: 628 Keats Court, Palo Alto, CA 94303
Citizen of: USA

SMALL ENTITY: YES

UNDEFEATABLE TRANSFORMATION FOR VIRTUAL MACHINE I/O OPERATIONS

BACKGROUND OF THE INVENTION

5 FIELD OF THE INVENTION

[0001] This invention relates very broadly to a computer system, and in particular to a system and related method of operation for controlling input and output (I/O) operations between at least one virtual machine and systems external to the virtual machine.

10 DESCRIPTION OF THE RELATED ART

[0002] Input/output (I/O) operations are among the most common carried out on modern computers. Some I/O operations are required simply for the efficient operation of the computer within the limits of the existing hardware and are performed substantially without user intervention, sometimes not as the result of any user action at all. For example, an operating system (OS) may need to swap data between a fast, volatile RAM memory and a hard disk simply because there is not enough RAM available for some immediate need.

[0003] Other I/O operations arise because of some user action. For example, when a user wants to send e-mail via a network, he may click on a "Send" icon in an e-mail application. Several I/O operations must take place in order to carry out this command. First, the displayed icon itself represents an output operation from the browser application, via the OS, a display driver, and a graphics card, to the monitor that the user is looking at. Second, sensing the position of the displayed cursor and the click of the input device (such as a button on a standard mouse) both involve inputs that are routed to the application via the appropriate driver in the OS. Third, in order to complete the "Send" operation, the OS must transfer the data (and associated transmission information) as an output operation to the driver of some device such as a modem or equivalent network connector, which then coordinates the actual transfer of the data out to the network.

FOIA b 7 - Excluded

[0004] Each I/O operation is a potential security weakness. A very common example of this would be destructive computer viruses: If users never needed their computers to be connected to any other computers, and never loaded data that came from any other source, and all their native software and hardware were certain not to have any viruses to begin with, then they would never need to worry about viruses. This would of course be the equivalent of living one's entire life in a sterile environment without any chance of contamination from the outside world. This is also of course impractical in all but very specialized cases.

[0005] Other problems of "security" have nothing to do with destruction of valuable data or damage to the computer's software itself. As many concerned parents have come to realize, the Internet makes readily available to the eyes and ears of curious children much material that even the most jaded adult would be far better off not seeing or hearing. As a result, many software packages are now available that attempt to filter the input from the Internet to the users' display screen and/or loudspeakers. Because such commercial network filters are applications loaded into the computer like any other, they can be defeated, often more easily by the computer-savvy children they are meant to protect than by the concerned parents.

[0006] Yet another frequent security concern involves the privacy of data transferred outside the computer (for example, over a network), or even within a single computer that has multiple users. The most common method for dealing with this is encryption. In conventional systems, some software package is installed in the computer that encrypts and decrypts the data before and after I/O operations. Once again, if the encryption routine is implemented purely as software that is at the same functional level as other applications, then it can be defeated or rendered non-functional. To minimize this weakness, some systems include specialized encryption hardware that resists software attack. Although generally more secure, the strength of this solution is also its weakness -- the user must buy and install specialized hardware.

[0007] Other issues of I/O control have nothing to do with security as such. Companies that maintain frequently accessed general web sites, such as those for major news organizations, web-based e-mail systems, etc., agree to insert into their

1
2
3
4
5 sites banners or other forms of displayed advertising that users see whenever they
view the site. Of course, users are generally as keen to be rid of such irritating
advertising clutter as the advertisers are to have it displayed. Accordingly, filtering
software exists that tries to analyze the content of each particular web page and
eliminate the advertisements.

6
7
8
9
10 **[0008]** From the perspective of the user, the ability to eliminate advertising banners
may be a blessing, but in this and other cases it may be undesirable to allow the user to
defeat such a display. For example, employers who have an internal company network
may want to make sure that the company logo always appears on any display, or that a
11 banner or tag such as "Company Confidential" always displays when an employee is
viewing a screen that is displaying information from certain sensitive files such as the
information in personnel files. Again, as in previous examples, any software application
one might install to prevent the user from defeating the display itself could be defeated
12 by a sufficiently skillful user, that is, one who knows enough about the system software,
hardware and the application itself. Note that even supposedly "impenetrable"
13 software has been defeated ("hacked") in recent years, often by users accessing the
respective systems remotely, and often by "experts" who are so young that their
parents have to drive them to their criminal trials.

14
15
16
17
18
19
20 **SUMMARY OF THE INVENTION**

21
22
23
24
25 **[0009]** The invention provides a method for controlling input/output (I/O) operations of
a user's computer, as well as the hardware and software architecture to carry out the
method. According to the invention, the user's computer is implemented as a virtual
machine (VM). A virtual machine monitor (VMM) is also included as a VM-transparent
26 interface between the VM and a physical computer system, which includes at least one
device. The VMM includes software modules for sensing a request for an I/O operation
between the VM and the device and for performing a predetermined transformation
(which may optionally even be time-varying) of I/O data passing between the VM and
the device. The I/O data may be of any type, including video (display) or audio data,
27 network packets, data being written to or read from a disk, etc. The transformation of
28
29
30

the I/O data is thereby made undefeatable by any user action via the VM.

[0010] In one embodiment of the invention, the device is a display and the I/O data is VM display data output from the VM and intended for display. In this case, the predetermined transformation is a replacement of at least a portion of the VM display data with non-defeatable display data stored external to the VM but accessible to the VMM, for example, in a display map/buffer located within the VMM. The data in the map/buffer is then used as the basis for the display such as that is visible to the user; the display will include the non-defeatable display data as an overlay.

[0011] Not all I/O data to or from the VM may always need to be transformed. Rather, the I/O data may be filtered with respect to at least one predetermined filtering condition and will be transformed only when the filtering condition is met. Many different filtering conditions may be used by the invention, depending on the type of I/O data and the device it is being sent to or input from. For example, the filtering condition may be that the I/O data includes at least one predetermined restricted term, or is from a predetermined restricted source.

[0012] According to one aspect of the invention, the I/O data may include image data. In this case, the step of filtering the I/O data comprises detecting the presence of a representation of a target image within the image data and the predetermined transformation is substitution of a representation of a replacement image in place of the representation of the target image. The I/O data may even be in a non-character image format, in which the target image is a representation of a restricted character string rather than an actual character string. In this case, character recognition may be used to convert the image to text, which may then be filtered by conventional parsing of the recognized text.

[0013] Another example of a filtering condition is the presence in the I/O data of a copy protection indication. It is also possible to use the invention to insert a source indication associated with the VM into the I/O data.

[0014] In other applications of the invention, the device is a network connection device. Again, many different types of transformations may be applied to I/O data between the VM and the network connection device. For example, the I/O data

transfer, on either transmission or reception or both, may be subjected to bandwidth limiting, or may be delayed, or (on transmission) redirected to a different address.

[0015] Yet another I/O transformation that the invention may provide is encryption/decryption. For example, data that the VM writes to a physical memory disk can be encrypted, and then decrypted upon read operations to the VM.

[0016] In VM's that support a plurality of I/O modes (such as audio and video, or from a network to video, etc.) the invention may filter the I/O data with respect to a criterion relating to one of the modes but transform data expressed in a form corresponding to a different one of the I/O modes.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] Figure 1 is a block diagram that illustrates the main hardware and software components of the invention, including an I/O control module within a virtual machine monitor on which a virtual machine runs.

[0018] Figure 2 is a block diagram that shows the main software components of the virtual machine monitor used in the preferred embodiment of the invention.

[0019] Figure 3 illustrates one example of how the invention may be used, in this case to generate a display overlay that cannot be removed by the user.

DETAILED DESCRIPTION

[0020] The main idea of this invention is a unique way to control the flow of information between a user's computer and all systems external to that computer. In the following description, the main hardware and software components of the invention are described. Then, the general method of operation according to the invention is discussed. Finally, several examples are given of how the invention can be used in different applications.

MAIN SYSTEM COMPONENTS

[0021] As Figure 1 shows, the main components of a computer according to the invention include an underlying hardware platform 100, host system software 200, and at least one virtual computer (also commonly referred to as a "virtual machine") 300, ..., 300n. In the preferred embodiment of the invention, a virtual machine monitor (VMM) 400, ..., 400n is also included for each virtual computer. These main components are discussed below.

System hardware

[0022] The system hardware 100 includes a central processor CPU 110, which may be a single processor, or two or more cooperating processors in a known multiprocessor arrangement. As in other known systems, the hardware includes, or is connected to, conventional registers, interrupt-handling circuitry, and other components as needed for normal operation of the system. These components are not shown or discussed further here since they are so well understood in the art of computer design. As in most computers, two different types of data storage devices are preferably provided: A relative fast system memory 112 (the machine memory), typically implemented using any of the various RAM technologies as part of the hardware platform 100, and a much slower but usually higher-capacity storage device such as one or more memory disks 114.

[0023] The system hardware will also include any components needed to enable connection to and data transfer with any external systems such as a network 700. For

example, Figure 1 shows the hardware platform as including a conventional network interface card (NIC), which is or acts as a modem or analogous network connector.

System software

5 **[0024]** As in other known computers, the system according to the invention includes system software, which typically comprises an operating system (OS) 220. The OS will include drivers DRV 222 as needed for controlling, communicating with and scheduling various installed devices 500. These devices 500 will typically also include various conventional input/output devices such as a display monitor, a keyboard, some pointing
10 or cursor control device such as a mouse, trackball, touchpad, etc., and possibly such peripherals as a printer, video and/or graphics card. No particular peripheral devices are required by the invention although some, such as a printer or display, will normally be included at least for administrative purposes. Note that the disk 114 itself may be considered to be an installed device in that read and write access to it is typically
15 scheduled and handled by the OS substantially as for any other device. This is not required by the invention, however.

[0025] The OS also includes any other conventional software components (not shown) necessary to handle resource (such as I/O device, memory, etc.) allocation and scheduling for installed applications 600. One software component found in many
20 operating systems is an application program interface (API); if such an interface is required for any application 600 loaded on the OS 220, then one will of course be included in any implementation of the invention. Because this invention does not presuppose any particular OS 220, and because the characteristics and functions of operating systems are so well known, the OS 220 is not discussed in greater detail
25 below, except insofar as it cooperates with the components of the system unique to the invention.

Virtual machines

[0026] In the preferred embodiment of the invention, at least one virtual computer
30 300, ..., 300n, usually referred to in computer science as a "virtual machine" (VM), is

installed to run on the host. As is well known in the art, a VM is a software abstraction -- a "virtualization" -- of an actual physical computer system. As such, each VM 300 will typically include a virtual CPU 310 (VPROC), a virtual operating system 320 (VOS) (which may, but need not, simply be a copy of a conventional, commodity OS), virtual system memory 312 (VMEM), a virtual disk 314 (VDISK), virtual peripheral devices 350 (VDEVICES) and drivers 322 (VDRIVERS) for handling the virtual devices 350, all of which are implemented in software to emulate the corresponding components of an actual computer. Although the key components of only one VM 300 are illustrated in Figure 1 and discussed below, the structure of any other VM's will be essentially identical.

[0027] Of course, most computers are intended to run various applications, and VM's are usually no exception. Consequently, by way of example, Figure 1 illustrates a group of applications 360 (which may be a single application) installed to run at user level on the VOS 320; any number of applications, including none at all, may be loaded for running on the VOS, limited only by the requirements of the VM. If the VM is properly designed, then the applications (or the user of the applications) will not "know" that they are not running directly on "real" hardware. Of course, all of the applications and the components of the virtual machine 300 are instructions and data stored in memory, just as any other software. The concept, design and operation of virtual machines are well known in the field of computer science.

Virtual machine monitors

[0028] Some interface is usually required between a VM and the underlying "real" OS 220 (real in the sense of being either the native OS of the underlying physical computer, or the OS or other system-level software that handles actual I/O operations, takes faults and interrupts, etc.) and hardware, which are responsible for actually executing VM-issued instructions and transferring data to and from the actual, physical memory and storage devices 112, 114. This interface is often referred to as a virtual machine monitor (VMM). A VMM is usually a thin piece of software that runs directly on top of a host, or directly on the hardware, and virtualizes all the resources of the

machine. The interface exported to the respective VM is the same as the hardware interface of the machine, or at least of *some* predefined hardware platform, so that the virtual OS cannot determine the presence of the VMM. The VMM also usually tracks and either forwards (to the OS 220) or itself schedules and handles all requests by its VM for machine resources as well as various faults and interrupts. The general features of VMM's are known in the art and are therefore not discussed in detail here.

[0029] In Figure 1, VMM's 400, ..., 400n, are shown, acting as interfaces for their respective attached VM's 300, ..., 300n. It would also be possible to include each VMM as part of its respective VM, that is, in each virtual system. Moreover, it would also be possible to use a single VMM to act as the interface to all VM's, although it will in many cases be more difficult to switch between the different contexts of the various VM's (for example, if different VM's use different virtual operating systems) than it is simply to include a separate VMM for each VM. The important point is simply that some well-defined, known interface should be provided between each virtual system 300, ..., 300n and the underlying system hardware 100 and software 220.

[0030] In some conventional systems, VMM's will run directly on the underlying system hardware 100, and will thus act as the "real" operating system for its associated VM. In other systems, the OS 220 is interposed as a software layer between VMM's and the hardware; still other arrangements are possible. This invention works with all such configurations, the only requirement being that I/O requests by any VM, and the returned results of the requests, should be able to be tracked and intercepted by some software component that performs the functions of the VMM described below.

[0031] In the following description of the invention, merely for the sake of simplicity, only one VM/VMM pair is discussed. The discussion applies equally, however, to all such VM/VMM pairs that may be included in any given implementation of the invention.

Interrupt handling and device emulation

[0032] Both hardware and software components can issue I/O requests via the operating system (or whatever software handles such requests). Depending on the type of request and why it was issued, the requesting component can either enter a wait state or it can continue with other tasks until the request is fulfilled. In both cases, however, the requesting component needs to be notified when the I/O request has been completed. This is typically handled using interrupts: When the requested I/O has completed, the device that carried it out issues an interrupt, which is usually a signal to the processor. An interrupt handler, conventionally within the OS 220, then senses the interrupt, notifies the requesting component, and makes the results of the I/O operation available to that requesting component using a predetermined protocol.

[0033] In the context of this invention, whenever an application 360 or device 350 running on or via the VOS 320, or the VOS 320 itself, issues an I/O request, this request is intercepted and handled by the VMM. This is known in the art of virtual machine technology; indeed, this is a primary function of a VMM, almost by definition.

[0034] As Figure 1 illustrates, the VMM 300 preferably includes an interrupt handler 430, which is a software component well known to computer scientists. Assume that the VM (in particular, the VOS 320) has issued a request for a page of memory from disk. The request will be addressed to the virtual disk VDISK 314, but the actual data will be stored on the physical disk 114 -- the VMM will, as in conventional systems, include an appropriate memory map and memory management unit 460 to handle translations between virtual and physical memory spaces, to perform memory tracing (if needed) on the memory allocated to the VM, etc. The VMM will then pass the request for the corresponding page of physical memory either to the OS 220 or will handle it itself in systems where the VMM itself directly manages the resource (here, the physical disk). Once the requested page has been transferred to a designated page in the system memory 112, an interrupt is raised by the disk driver to signal completion of the transfer.

[0035] The interrupt handler 430 in the VMM then intercepts this interrupt and maps the page of system memory to a corresponding page in virtual memory. The VMM then

signals completion of the transfer to the VOS, in effect simulating another "interrupt" to the VOS. The VM then continues execution as normal. From the perspective of the VM, that is, of any user of any application running in the VM, all that will have happened is that a normal transfer from disk to memory will have occurred. In reality, however, the virtual transfer will correspond to a physical transfer, with the operations performed by the VMM (here, memory mapping and interrupt handling) remaining transparent to the VM.

[0036] In order to provide more generality, the VMM according to the invention may also include device emulation modules 450. As the name implies, these modules are software interfaces between the VMM and the VOS 320 that emulate the characteristics of respective physical devices.

I/O control

[0037] In this description, the operations controlled by the invention are referred to as "I/O". Note that the concepts of input and output depend on the perspective of two communicating components, so that even when all data flow controlled by the invention is *input* to the VM, it is still, simultaneously, an *output* from the device that completed the I/O operation requested by the VM; it would also be an output operation from the VMM, which forms the interface between the VM and the physical, non-virtual system. The same applies to applications of the invention in which the only data transformed is from the VM to some external device -- such output from the VM would still be input to the external device. The term "I/O" is therefore not intended to limit the invention to use in transforming data that passes in only one direction, that is, either to or from the VM.

[0038] One module within the VMM according to the invention that is not found in the prior art is an I/O control module 470. Several examples of the function of this module are given below. The function can, however, be described generally as follows: Whenever the VM requests any of one or more predetermined I/O operations, then before the results are actually made available to the VM, they are transformed by the I/O control module 470, more specifically by a transformation module 474 and possibly

only when a filtering module 478 detects that the I/O data satisfies some transformation-triggering criterion or condition (these modules are discussed below).

[0039] The invention is able to carry out many different types of transformations, on many different types of data. For example, where the I/O involves some visible presentation of the data, transformations might encompass deletions, insertions, and modifications of the presentation. Transformations may also be temporal: Where the I/O involves transfer of data between a VM and the network (or some other device that streams data), for example, the transformation applied may involve introducing transmission or reception delay or bandwidth limiting. Many other examples of transformations are given below.

[0040] The operation of the I/O control module can be activated within the VMM, for example, when the interrupt handler 430 signals that an I/O request has been completed. The transformation applied by the module 430 may also be "triggered" when the user takes certain predetermined actions, or when the data meets certain predetermined transformation or filtering criteria. This is explained further below.

[0041] The I/O control module 470 may transform not only data input to the VM as the result of a VM-generated request, but also data that is to be transferred (output) from the VM to a device or software routine that is external to the VM. The I/O control module 470 may also operate in a "streaming" mode, in which data passing to and/or from the VM via the VMM is transformed serially, such that the I/O control module does not need to wait for the entire results of the transfer to be available at once, but rather can operate on smaller I/O units, for example, pages, files, or even individual words or bytes or bits (for example, parity bits).

[0042] Figure 2 shows the structure of the I/O control module 470 in greater detail. In Figure 2, the OS 220 and various hardware components are not shown merely for the sake of simplicity -- their presence can be assumed, and they will perform their normal functions. Figure 2 also shows various components of the I/O control module 470 that are used by the invention when the VM issues an instruction that involves I/O with some device 502 that is external to the VM. This device 502 may be part of the system as a whole, for example, the physical disk 114, or a device that is accessible only via a

network connection; the principles of operation of the invention are same in both cases.

[0043] The device 502 need not involve hardware; the only common concept in the context of this invention is that the device is some structure that is not within the VM, but with which the VM communicates via the VMM.

5 **[0044]** Data (which here can also be taken to include executable code, since it too is typically transferred in the same way as non-executable data) that is to be transferred to (or from) the VM as a result of an I/O request is stored in a memory space or module 472 referred to here as the "map" or "buffer" or "mask" depending on what type of information it is currently being used to store. For example, as is explained below in
10 greater detail, this module may be a bit map used for a display, or it may be a buffer space where data transferred serially via a network is first assembled before it is transformed.

[0045] The data, or a sufficient part of it, that is passed to the VMM for forwarding to the requesting VM is preferably examined by a filter module 478, which applies one or
15 more predetermined filtering conditions to decide whether the data is to be transformed or not; the filter 478 may be deleted if I/O data for the VM is *always* to be transformed -- examples of when one might prefer this are given below.

[0046] An administrative system 650 (an application or even another computer system, which may even be another VM, and is under the control of a human
20 administrator) may also be connected to the VMM either directly or, in most cases, via the OS 220 and whatever software is included (for example, a driver 222) used for communication between the OS and the VMM, in order to set the parameters of the I/O transformation applied by the I/O control module 470. Depending on the type of transformation to be applied, these parameters may include the filtering condition(s)
25 mentioned above. Figure 2 also illustrates a user 800 who is operating the VM.

Examples of non-defeatable I/O transformations

[0047] In order to make the operation and structure of the invention more understandable, consider the following example. This example is also illustrated in Figure 3, in which, for the sake of simplicity, other components of the system as a whole (which are shown in Figures 1 or 2 or both) that are not needed for understanding of the example have been omitted, but can be assumed to be present and performing their normal functions. In Figure 3, the I/O control module 470 is shown having as its main software components not only the display map 472, but also a transformation module 474 and a filter module 478.

[0048] Assume that the user is working with a standard spreadsheet application and wants to open a file containing sales data, which he wants to view on the screen 910 of a physical display monitor 900. Because the user is interacting directly only with the virtual computer 300, any data for display will be maintained in an allocated portion of the virtual memory 312.

[0049] The request for the file -- represented graphically in Figure 3 as a bar chart 480 stored within the physical disk 114 -- will pass as normal from the application to the VOS 320, which will then retrieve the corresponding disk blocks in its normal manner from the virtual disk 314. The retrieval of the data from the virtual disk is intercepted by the VMM, which performs the I/O operation on behalf of the VOS, and possibly transforms the data that is being read or written to during the I/O operation.

[0050] As in conventional virtual systems, the file (that is, the disk blocks associated with the requested file) may then be transferred as is into virtual memory for use by the application that requested it. A corresponding "real" display map -- the one that is accessed by the display driver located in the OS 220 -- is, however, maintained in the VMM, or in some component that is accessible to the VMM, preferably in the I/O map/buffer space 472.

[0051] In this example of the invention, assume the file (associated disk blocks) is designated as being restricted, which may be determined by comparing the corresponding file, directory or drive name, or disk block identifiers, with a list maintained in the filter module. The transformation module will then generate and

apply to the display map 472 an overlay in the form of a visual indication 920 such as a banner or tag that indicates access (successful or attempted) to the confidential file. In Figure 3, the banner reads "CONFIDENTIAL FILE ACCESS." This can be done simply by writing the corresponding bits (or ASCII bytes, or whatever minimum display units
 5 are used) into the display map 472. As in conventional display data, color and effect information may also be part of the overlay, so that the banner could be, for example, always in red, or could blink, or could be some other dynamic, time-varying transformation such as an animated overlay, etc.

[0052] Filtering could also be performed based on the higher-level name of the file(s) to be transferred, assuming these names are visible to the VMM. In the case of a file
 10 read from a network server, for example, the filtering module could observe the NFS or SMB protocol messages, which may include file and server names, and direct the transformation module to perform the predetermined transformation on the I/O data any time the name of a requested file is on a target list, or not on a permitted list.

[0053] The banner or tag (such as "CONFIDENTIAL," or a company logo) may be placed within the graphical display of the virtual machine, as Figure 3 illustrates. If the resolution of the display of the virtual machine is less than the capabilities of the real frame buffer, however, the transformation module could place the banner or tag outside the normal display of the virtual machine, for example, either below, above, on the right
 15 or the left of the virtual display.

[0054] The result of this transformation operation is that the banner or tag or other visible indication 920 will appear on the physical display regardless of what the user does -- because the user's actions are restricted to the VM, he has no way to influence the actual display map 472 used to generate the physical display. In other words, the
 20 user 800 cannot defeat (prevent or undo or in any way affect) the transformation of the requested I/O operation even if he somehow were able to gain direct control of the VOS 320 and all other components of the VM: untransformed data is simply never available for access by the physical display driver because the VMM is transparent to the VM.

[0055] In order to make the transformation completely transparent to the VM, the
 25 original, untransformed display data may be allowed to be stored as is in the VM frame

(display, audio, etc.) buffer. It would also be possible, however, to transform even this data before storing it in the VM frame buffer. If this is done, and especially if no copy of the "raw" data is stored within the VMM or in any other buffer, then there will no copy of the raw data stored anywhere in the entire system.

5 **[0056]** In Figure 3, the display 900 is shown as rendering the bar chart, that is, the requested file, along with the banner 920. This is not necessary. Instead, the banner could be displayed regardless of the current display merely to indicate that the user has, or has attempted to, input a restricted file. The requested information itself could be displayed or not depending on the user's access privileges.

10 **[0057]** Moreover, just because data may be intended primarily for display, it is not necessary that the transformation be limited to the display data. Instead, or in addition, the banner (or other notice), could be inserted into, for example, buffered print data, so that it would appear on any print-out. In this case, the transformation could be applied to data used for presentation via two different external devices -- a display and a printer.

15 **[0058]** It would also be possible for the overlay to be translucent, letting through the original display data with different levels of visibility. For example, if the original text is mostly visible (very translucent overlay), then the overlay will take more the appearance of a watermark. Allowing the overlay to mostly obscure the original display data, however, will have the effect of calling the attention of the user to the transformation, while still letting him see what caused it. The way in which a translucent overlay is created is well known in the art of computer graphics and involves "interspersing" in the display map the bits of the overlay with those of the original data.

20 **[0059]** Using the invention to combine original I/O content with a translucent overlay is not restricted to on-screen displays. For example, as is mentioned briefly above,
25 whenever a user prints a document, the transformation module may be set to generate a watermark and incorporate it into the data in the printer's buffer. This watermark may be made very weak and difficult to detect. If the watermark is applied only to data to be printed that has predetermined characteristics, or comes from specific sources, then this watermark would be able to function as a non-defeatable indication of copying.
30 Instead of a translucent overlay, the non-defeatable indication could also be some

predetermined but hard-to-see mark placed somewhere in the print frame. These techniques are sometimes referred to as "copyright traps" since they provide all but irrefutable evidence that the data being printed was not generated originally by the VM user. The advantage of using the invention to accomplish this, though, is that even if the VM user knew of the nature of the trap, he would still not be able to prevent it from being generated and printed.

[0060] The operation of inserting hard-to-detect markings within the I/O data may also be extended to data created by (if this can be determined) or at least passed from the VM. For example, the transformation module could insert some source indication, such as a predetermined bit pattern or alteration into the I/O data sent from the VM via the VMM for forwarding to a device such as a printer. This would then serve as irrefutable evidence that the VM had at one time accessed the data. Care must be taken, of course to avoid introducing check-sum or parity errors in the I/O data in cases where such techniques are used for error detection by the recipient device.

[0061] The filtering module could also be programmed to detect the presence of such a copy-protection indication within I/O data. Such indications are already being included, for example, as part of the protocol of digital audio and video transmission formats. If the filter module detects the presence of the copy-protection indication in the I/O data, then the transformation applied by the VMM could be, for example, the forwarding of an error message to the VM instead of the expected data, including replacement of the entire set of I/O data by the error message.

[0062] This concept of undefeatable copyright traps can also be extended to other modes, such as audio. Thus, for example, into the stream of every audio file downloaded from the network with some predetermined file extension (such as .mp3), or into the buffered downloaded data, the transformation module could insert or overlay some predetermined bit pattern (audio mark) that would identify the file as having been obtained via downloading. The filter module could in this example also be provided with a list of approved domains (for example, ones from which the owner has bought a license) from which, say, mp3 files may be downloaded without triggering the generation of the mark.

Filtering

[0063] It is of course possible that only certain data to be displayed needs to be marked with the non-defeatable banner 920. In these cases, the VMM could, for example, use conventional memory tracing to mark those portions of disk storage that contain company confidential material. Any transfer of data from a marked portion could then be used as a trigger for the transformation module 474. In this case, the filter condition would be the requested transfer of any such marked data to the display map.

[0064] It is not necessary, however, to apply any filtering at all; rather, *all* displays generated by the VM could be transformed so as to include a non--defeatable banner 920. For example, the transformation could be an overlay of, for example, a company's logo or an advertisement. In such case, the filtering condition would be simply that any data is present in the display map 472 used by the physical display 900.

[0065] To consider another example of a use of the invention, assume that the external device 502 is a network interface connected to the Internet. The Internet would thus replace the disk 114 in Figure 3, and the application the user 800 is using will typically be a browser. Here, display transformation can be either conditional or unconditional, or both, depending on content.

[0066] In the simplest case, *all* displays generated by the browser could be transformed by setting the transformation module 474 to add as a display overlay a non-defeatable banner advertisement. One reason the user might even willingly tolerate this would be if, for example, the manufacturer of the system as a whole were to sell it to the user at an attractively low price, but only on the condition that he accept the transformed displays.

[0067] Note that because a VM, by definition, acts precisely as a "real" computer, it might not even be necessary for the user ever to know that the computer on which he has loaded his applications is entirely virtual. This transparent operation is, in fact, more common than one might suppose: Few users of the widespread Microsoft
5 Windows 98 operating system realize, for example, that some of the older applications they may choose to run are actually executed using a virtual machine intended to accommodate the appropriate legacy operating system or the outdated 80386 Intel processor.

[0068] In the examples given above, a visible overlay is added to the display map,
10 either always or only if some condition is met. It would also be possible for the "overlay" to consist wholly or partially of some "null" value, whereby one or more portions of the display are masked out, that is, simply not displayed. This, too, could be either conditional or unconditional. For example, assume that one determines that banner advertisements are typically located in a 3 cm horizontal strip whose top edge is 4 cm from the top of the screen displayed by a web browser, just under the toolbars. The transformation module 474 would then mask out the 3 cm strip, for example by setting the corresponding bits in the display map to "black."

[0069] Banner advertisements that include text are usually generated from code, for example, in HTML, XML, or some equivalent, that can be parsed. This means that the filtering module 478, using well-known programming techniques, could examine the
20 HTML (for example) code in order to detect the presence of predetermined text and either delete or replace only that text. For example, assume that the "raw" current HTML code includes, in a banner, the term "Brand X." The transformation module could then either delete "Brand X" and allow the rest of the banner to display, or it could
25 replace this text with something else, for example, "Brand Y", or a warning against Brand X.

[0070] The invention is not restricted to filtering and transforming displays based on parsed text in underlying source code. It would, for example, also be possible for the filtering module to operate directly on raw pixel data, that is, the bit (or vector, etc.) data
30 in the display map 472, and thereby to transform images as well. Assume, for

example, that the administrator of the system (or whoever programs the filter and transformation modules) wishes to prevent users from viewing certain target images, such as a predetermined logo or symbol (such as a swastika). The filtering module 478 would then include a representation or characterization of the target image. The pixel pattern stored in the display map would then be scanned and analyzed using any conventional algorithm to determine the presence and location of the target image. If the filtering module detects the presence of the target image, then the transformation module replaces it with some other predetermined image, which may be a simple pattern of black pixels or even a complex replacement logo or symbol.

[0071] Note that text may in many cases also be represented as a non-ASCII image, for example, where a document has been downloaded in an image format such as pdf, gif, or jpg. In this case, the presence of target text could be determined by applying a known optical character recognition (OCR) algorithm to the pixel set in the display map and parsing (if necessary) the converted image.

[0072] Regardless of how the presence of a target image is detected, once found, the transformation module would then be invoked to transform the display in any appropriate way. The target image could be blacked out, for example, or replaced with some other image, or covered with a banner. The I/O control module 470 could then also (or instead) send a message, for example, to the administrator or via the network to some other party, giving notice that the user had downloaded or attempted to view or actually viewed the target image.

[0073] The pixel (display) map may also be filtered and transformed based on target characteristics other than the content of some specified image or text. One of the many alternative characteristics that may be used to trigger transformation would be the color of some portion of the display. If, for example, a user suffers from the common red-green color blindness, then the filter and transformation modules could be set to identify all on-screen display portions that are to be displayed with one of the hard-to-see colors. Colors may, for example, be identified by their red-green-blue (RGB) or cyan-magenta-yellow (CMY) percentages, Pantone numbers, etc. If both red and green are found, or found too near each other, then the color one or both of the corresponding

display portions could be adjusted for greater visibility, for example, by replacing them with a different color.

[0074] Another example of a transformation that has nothing to do with any prohibited terms or images, but rather simply with a characteristic of text, could be a change in the font or size of text. This could be determined either by parsing the HTML/XML code (with the filter module acting as a network filter) or by using OCR on the display map. Text that is in a font too small or too unclear for the user to see clearly could then be automatically converted for display into a more desirable font, assuming there is enough room on the display.

[0075] It is not necessary for the display map 472 to be wholly or even partially a bit-mapped or vector display representation; rather, in some applications, it may act as a simple buffer that stores symbolic (for example, ASCII) information that is converted for display by some other module, for example the display driver . This symbolic information can then be parsed and transformed as described above.

[0076] The filter module is not restricted to operating on data used directly for display, such as a bit map. Instead, acting as a network filter, it may instead operate on some other stream of code or data that is used to generate the display. One example of this would be an HTML or XML stream sent to a browser. In this case, the contents of the stream can be parsed by conventional code within the filter module 478. If the filter module detects one or more words from a pre-stored list (entered, for example, by the administrator), then the transformation module could remove from the stream the word(s), or the sentence or paragraph containing the prohibited words, or all text within some predetermined distance (for example, 2 cm above and below) of the word(s). The invention would in this implementation be performing the same function as a conventional network filter, for example, to protect children from pornography. The difference, however, is that there is no way for even the most computer-savvy teenager to override the filtering operation, because he is isolated within a virtual computing environment and the filtering and transformation code is not within the VM.

[0077] Transformation could also be based on the I/O source instead of, or in addition to, the I/O content. In this case, the filtering module 478 could contain a list of restricted sources (for example, e-mail addresses and/or URL's, or disk drives or memory partitions), or a list of permissible sources. Any information coming from a source on the restricted list, or not on the permissible list, could be "transformed," for example, by having the entire display "blacked out", along with a banner notice as to why.

[0078] This embodiment may thus also be used to perform the function of an existing network content filter. Unlike conventional filters, however, the invention allows for a feature that is not available in the prior art, namely, the non-defeatable transformation. Instead, a non-defeatable message could be inserted into the display map alerting the user to the fact that input from an unauthorized source was attempted. Moreover, using the invention, there is no need to install any special filtering software within the VM itself; indeed, the invention can operate without any modification to the VM at all.

[0079] Yet another transformation criterion may be not the source but rather the destination of network traffic from the VM. For example, if the VM attempts to send data to a network location (for example, a URL) on a restricted list maintained within the filtering module, then the transformation module could redirect the transmission to some other destination simply by inserting the substitute address into the transmission at the proper place, which can be determined by and will depend on what network protocol is used. One example of when this might be useful is in systems where high-speed network routers are implemented in parallel to a standard network. The invention would then automatically redirect VM network content to the high-speed network routing system.

[0080] The invention is not restricted to transforming the contents of the display map (or other buffered I/O information) based on what is coming in to the VM, but may also be used to transform information being output from the VM in a way that cannot be overridden by the user. Assume, for example, that the user attempts to send an e-mail message to a recipient either not on an authorized list (for example, on an authorized server), or to a recipient on a "blocked" or "forbidden" list, or with a message containing a forbidden term. Using the invention, the filter module would examine the address of

the intended recipient and/or the content of the message, and if it is not permitted, could either "transform" the address (or terms) by deleting it completely, thus blocking the "send" operation altogether, or it could redirect the address, with or without transforming the content as well, to some other recipient, for example, a parent's computer, or a security department.

[0081] This possibility is not restricted to the elimination of pornography. Rather, as another example, this aspect of the invention could be used to prevent inadvertent display or disclosure of the secret internal name of a project -- even if the user typed the name, it would not be displayed, or printed, or output to the network.

[0082] The VMM, via the transformation module, could also prevent all files or messages containing specified restricted terms, or coming from predetermined sites or files or disk drives, from being written to any other memory device or region, or printed. In this case, after filtering, the applied transformation is a marking of the data in such a way that the VMM will not allow it to be rewritten or even printed. Thus, e-mail messages or data files could be converted into a "FOR YOUR EYES ONLY" format (which itself can be displayed as a tag for the user to see) that cannot be defeated by any user action because it is taking place outside of the virtualized environment to which the user is restricted. As in other cases, this can be done without installing any specialized software in the VM, and without modifying any software already running in the VM.

[0083] The invention may also be used to place a non-defeatable notice on a display, or in a printer buffer, etc., whenever predetermined filtering conditions are met, such that the notice can be turned off or otherwise removed only by the action of someone other than the user of the VM, for example, the administrator. Just a few of the examples of when this feature might be advantageous include when the user has attempted to access a forbidden site (for example, outside an employer's domain), or has attempted to access a disk file that he is unauthorized to access (for example, the personal information about company employees, or information with a security classification above that for which the user is cleared), or has attempted to download or read in data from a site or file that contains a prohibited term, or has tried to print a file

containing a prohibited term, or has tried without authorization to copy a file from the virtual disk onto a portable memory device such as a floppy disk.

[0084] In all of these cases, once they have been detected by the filtering module, the transformation module can place a tag or similar notice on the display associated with the user's VM, or in a printer buffer of a printer, etc., as appropriate, that cannot be removed by any action on the part of the user. The type of notice will depend on the situation, of course, and could include such phrases as "UNAUTHORIZED ACCESS ATTEMPT" or "UNAUTHORIZED OPERATION" or "SECURITY VIOLATION", etc. Alternatively, a banner such as "SECURE" could be displayed on the screen *until* the user attempts some unauthorized I/O operation, at which time the transformation module removes the banner from the display.

[0085] If the filter module 478 includes conventional virus-detection software, then the filtering condition that triggers a display transformation could be that a virus is detected in some incoming or even outgoing data stream or file. The transformation module can then place a notice such as "VIRUS WARNING" on the display, and the VMM can automatically isolate the corresponding file in the buffer 472 until the user or administrator takes some predetermined measures to prevent infection. Of course, existing virus-protection software also displays a notice concerning detected viruses. Unlike this conventional software, however, the invention does not require the software to be loaded into the computer (here, the VM) that requested infected I/O data. Moreover, the invention can automatically check *outgoing* files, so that the VM would not be able to spread viruses further regardless of what the user does or neglects to do.

[0086] The invention can thus be used to implement a firewall that requires no specific hardware support, yet completely isolates every VM from every system outside the respective VM itself other than its associated VMM. This firewall is, moreover, completely unalterable and impenetrable by the user, since it lies outside the virtual computing environment to which the user is restricted.

[0087] The transformation module may also include several different non-defeatable banners or notices or filtering, in any combination. For example, files or network addresses (including internal) could be classified into different security levels, for

example, "CONFIDENTIAL", "SECRET", and "TOP SECRET", depending on their contents and/or which source the data comes from. As long as the VMM, the system hardware and software, and the administrator, have the highest security level, users with different clearance levels could then use the same system, but different VM/VMM pairs, and could be completely prevented from accessing data at a level above their clearances. The administrator can then set transformation modules 474 of each VMM to generate the appropriate security banner. In addition to preventing access, the VMM could also send a notice to the administrator any time a user attempts I/O of information for which he is not cleared.

[0088] One particularly advantageous form of transformation is encryption. In this case, before any data is transmitted from the VM to a physical device, whether internal (such as the disk 114) or external (such as via the network 700), it can first be encrypted by the transformation module using any known method. Note that the user himself need not even know the encryption routine or keys. Incoming data to the VM can then be decrypted and further filtered and transformed as desired before it is passed to the requesting VM. The encryption routine and keys can be set, for example, by the administrator, who has access to the VMM. If all members of a project worked only via respective VM/VMM pairs, with each VMM having the transformation module 474, then all (or certain predetermined) I/O to and from the users could be encrypted in a way that is completely transparent and inaccessible to the users.

[0089] One especially interesting use of transparent cryptological transformation according to the invention is when the physical device is the disk 114. If, as before, the transformation module encrypts all data that the VM writes to the disk (that is, the portion of the physical disk 114 on that contains the virtual disk 314), and decrypts all data read into virtual 312 memory by the VM from the disk, then no unencrypted data associated with that VM will ever reside on the disk 114. Indeed, one could physically remove the entire physical disk 114 but still not be able to access any VM-related data without knowing the proper keys.

[0090] One other example of the benefits of cryptological transformation using is where Ethernet packets are encrypted/decrypted during transfers between the VM and

a network. As before, any known encryption mechanism, such as IPSEC, may be applied by the transformation module. Once again, the encryption would be completely transparent to the VM user.

[0091] Similar to a cryptological transformation would be the transformation of *compression*. Accordingly, the transformation module could apply any known algorithm to compress all data to and from disk, or the network connector, and to decompress the data before its transfer to a requesting VM.

[0092] The transformation implemented by the invention need not involve any form of visible display. Encryption/decryption is one example of this. As one other example, the transformation module 474 may be programmed, using conventional techniques, to act as a network bandwidth limiter or network traffic shaper, for example, in order to allocate network bandwidth to different VM's. For example, each VMM could limit its respective VM to a predetermined maximum bandwidth such as 1MB/s; all I/O requests requiring more than 1MB/s of bandwidth could then be limited simply by having the transformation module 474 drop packets that exceed the 1MB/s bandwidth limit. Different VMM's could also cooperate in such as way as to dynamically adjust the bandwidth allocation for different VM's as a function of current availability.

[0093] The invention is well-suited for performing even other types of I/O transformations relating to I/O with a network. The function of network shaping is one example of a performance-based transformation. Another would be a transformation that deliberately imposed extra delay (latency) in an I/O operation. This type of transformation might be useful, for example, for testing systems. For example, the invention could then be used to simulate a lossy network connection by artificially dropping or causing bit errors in packets, or to simulate contention by delaying packets with some statistical distribution of delays. The delay may be introduced both in transmission of data from the VM to the network connection device and reception of data by the VM from the network connection device.

[0094] As another example, the invention may be used to transform *audio* data, in which case the external device 502 would be a sound card used to drive speakers. For example, assume that a video conferencing or telephone application is running on the

user's VM. The transformation module 474 could then encrypt the audio output data. Incoming audio data could then be transformed by being decrypted before forwarding it to the VM.

[0095] Instead of (or in addition to) encryption, the transformation module 474 could also automatically change the sound characteristics by applying standard filters, for example, to remove high-frequency components, to attenuate the volume of selected audio, and reduce the require transmission bandwidth. Similar to the manner in which the transformation module can insert visible banners into a display, it would also be possible for it to insert , or substitute, audio content into an audio buffer or stream. For example, the transformation module could insert audio advertisements, or reminder or warning messages.

[0096] The invention may also be used to perform multi-modal transformations, that is, a transformation of one I/O feature as a function of another. For example, the transformation in which one text string in a display is replaced by another; thus, a visible portion of the display (a visual mode) is replaced another portion of the display (still, the visual mode). As just one example of mixed modes, if the filter module detects the presence of some predetermined text string intended for display, then the transformation module could generate an audio signal in response. Thus, if "Brand X" is in the string, the transformation module could place in the buffer of a sound card prestored data that would cause an audible message such as "Have you ever tried Brand Y?" As another example, a parent's audible reprimand could be played whenever a child attempts to download files containing objectionable material (filtering of input) or to send e-mail or type text that contains objectionable terms (filtering of output).

[0097] In most of the examples of uses of the invention given above, the condition that triggers a transformation of I/O data is that the I/O data contains some characteristic, such as some prohibited text string, or comes from some prohibited source. The transformation-triggering filter condition might, however, just as well be the *absence* of some *required* text string (or file extension, etc.) or that the source is not

one of a set of approved sources. It will in many cases be easier to maintain a list of what is required rather than a list of what is prohibited.

[0098] In most of the examples of I/O transformation given above, "transformation" is used in the strict sense of the word, namely, a change of some data being transferred between the VM and some external device. It is to be understood, however, that in the context of the invention, it is not necessary for the VM to be the component that generates the I/O data, or even requests it; rather, an *insertion* of data into existing I/O initiated by some other component, for example, the administrative system 650 is also a kind of transformation.

[0099] For example, assume that one or more VM users are viewing respective display screens. The administrator, using conventional techniques, could then communicate with each (or selected) respective VMM by sending an appropriate command in order to direct its transformation module to insert into the corresponding VM's (user's) display map (or other I/O buffer) within the VMM some message, such as a banner. One example of such a message could be a notice of an impending shut-off of the system and a reminder to the users to save their work. Another example, useful in secure environments about to receive non-cleared visitors, could be that the administrator could direct each transformation module to black out all display screens except for a banner such as " NON-CLEARED VISITOR IN AREA -- PLEASE STAND BY," perhaps with some other display feature such as a flashing red symbol. This would absolutely prevent the uncleared visitor from being able to see, even inadvertently, sensitive on-screen information, even if a project member happened to be away from his work station or carelessly left a door open.

[0100] Note that this administrator-initiated display transformation (securing of the display) can be done for any or all users, but will require no additional hardware support and no changes to any software of the VM. Moreover, assuming that the user does not need to see the display while it is secured, the display transformation (here, being blacked out) could be done while the VM continues execution of whatever code is was executing before the transformation.

[0101] One general advantage of the invention that is mentioned above is that it is preferably implemented in a way that is completely transparent to the user. Another general advantage illustrated by many of the examples given above is that the invention can accomplish a very wide range of transformations on almost all types of I/O with no
5 need for specific hardware support at all, and with no need to install or modify any software in the VM or, for that matter, anywhere other than in the VMM.